

REMARKS

This amendment is being filed in response to the Office Action dated July 30, 2007. Various claims are amended as shown. New claims 60-62 are added herein. Claims 1-33 were previously canceled without prejudice. No new matter has been added. With this amendment, claims 34-62 are pending in the application.

I. Objection to specification

The present Office Action objected to the previously submitted amendments (which were submitted via a Substitute Specification) to the specification for containing information “that is identified as commercial site URLs or a form of hyperlinks ... It’s uncertain whether a standard browser or a future developed browser won’t execute the text URL as amended in the specification.”

To address this continuing objection, a Second Substitute Specification (clean copy), along with a Second Redline Substitute Specification showing the changes in redline format, are being submitted herewith. The Second Substitute Specification contains no new matter.

In the Second Substitute Specification, the format of the various items such as “www.gslb.com” and “www.foo.com” has been amended/alterd such that they are no longer browser-executable or browser-recognizable. Specifically, these items have been rewritten using ordinary and phonetic English-language text, including industry-accepted/recognizable terminology. For instance, the “www” occurrences in the previous specification have been re-written as --dub-dub-dub--, which is explained in www.dictionary.com as “[common] Spoken-only shorthand for the ‘www’ (double-u double-u double-u) in many web host names.” As another example, the “.com” occurrences in the previous specification have been re-written as --dotcom--, which is explained in www.dictionary.com as “from the pronunciation of .com, suffix of domain name in most commercial Internet addresses.” Therefore, “www.gslb.com” has been rewritten as --dub-dub-dub dot gslb dotcom-- in the Second Substitute Specification enclosed herewith. Corresponding similar changes to

ordinary/phonetic English-language text have been made to the other content of the specification that the present Office Action alleged to be “URLs” or “hyperlinks.”

In view of these amendments to the specification, it is kindly requested that the objection to the specification be withdrawn.

II. Supplemental Information Disclosure Statements (IDSs)

A supplemental IDS, along with an appropriate certification, was previously filed on January 7, 2008, so as to submit additional references for entry and consideration.

Another supplemental IDS, along with the requisite fee, is also being filed herewith, so as to submit still additional references for entry and consideration.

It is kindly requested that the Examiner provide initialed copies of both of these two supplemental IDSs along with the next communication, so as to confirm entry and consideration of the references listed therein.

III. Enablement rejections

The present Office Action rejected claims 43-46 and 47-50 under 35 U.S.C. § 112, first paragraph, for allegedly failing to comply with the enablement requirement. Specifically, the Office Action stated on page 3 (section 6) that a “machine-readable medium having instructions” is not disclosed in the specification.” Further, the Office Action stated that “a generic ‘machine-readable medium’ would include transmission line that is not statutory.”

The Office Action’s assertion that a “machine-readable medium having instructions is not disclosed in the specification” is respectfully traversed herein. For example, such a machine-readable medium is disclosed on page 10, line 24 through page 11, line 5 of the original specification as filed. In view of this disclosure, it is respectfully submitted that claims 43-46 meet enablement requirements.

To further meet enablement requirements, as well as statutory subject matter requirements under 35 U.S.C. § 101, claims 43-46 are amended as shown in a manner that is consistent with the original specification as filed. Specifically, “machine-readable medium” has

been amended to --storage medium-- having instructions stored thereon that are executable by a load balance switch or site switch.

IV. Discussion of the claims and cited reference

The present Office Action has rejected claims 34-59 under 35 U.S.C. § 102(b) as being anticipated by a white paper entitled “Server Load Balancing in Today’s Web-Enabled Enterprise” (hereinafter “the White Paper”). In view of the amendments above and the arguments below, it is kindly requested that the rejection of these claims on the basis of the White Paper be reconsidered and withdrawn.

A. Independent claim 34

Independent claim 34 as amended herein recites, *inter alia*, “providing said public virtual IP address from said site switch to at least one load balancing controller to enable said load balancing controller to update an address record to indicate said public virtual IP address as being associated with said site switch.” It is respectfully submitted that the present Office Action has not cites any passage of the White Paper that meets these specific limitations.

For example, page 4 (under the “SLB During Server Failure” subsection of the White Paper) provides the following description (emphasis ours) with regards to server load balancing (SLB) capability of a server load balancer, such as the server load balancer shown in page 5 of the White Paper that balances load amongst servers in a server farm:

“Further, SLBs also serve a role in returning the response to the client by translating the responding serves’ IP address, which the enterprise may not want to make public, to the publicly registered addresses.”

A few items can be discerned from the above-cited passage from the White Paper. First, the SLB device returns the response (having the “publicly registered addresses”) to the client, rather than providing public addresses (from the site switch) to a load balancing controller

as required in claim 34. Second, the present Office Action has not adequately cited any portion of this or other passages of the White Paper that meet the limitations of “said load balancing controller to update an address record to indicate said public virtual IP address as being associated with said site switch.”

Page 4 (section 9) of the present Office Action cited Figure 6 and the accompanying description of the White Paper in rejecting claim 34. It is respectfully submitted that the present Office Action also has not adequately identified any portion of these cited passages of the White Paper that disclose, teach, or suggest the limitations of claim 34 pertaining to “providing said public virtual IP address from said site switch to at least one load balancing controller to enable said load balancing controller to update an address record to indicate said public virtual IP address as being associated with said site switch.”

For example, page 6 of the White Paper shows a “Controller GSLB Switch” (CGS) and a “Site GSLB Switch.” Page 6 (the “DNS Lookup Process” subsection) of the White Paper provides the following description (emphasis ours):

“4. The CGS intercepts this packet [from the DNS server], selects the best IP address to send to the client in San Francisco ...

5. The CGS then sends this response back to the client DNS.

6. The Local DNS sends the information back to the client. The client will select the first address.

7. A connection attempt is made to the www.foundrynet.com site in Hong Kong for which the Authoritative DND provided the IP Address.”

From the above-quoted passage and accompanying figure on page 6 of the White Paper, a few items can be readily discerned. First, the CGS sends a response (having addresses of site switches) to the local DNS of the client (this communication from the CGS being represented by the arrow labeled as “5” in the figure). Second, the client is connected to the site switch in Hong Kong that was identified by the CGS as being the preferred site (this

communication from the client to the site switch being represented by the arrow labeled “7” in the figure).

The present Office Action has not cited any portion of the figure and accompanying description on page 6 of the White Paper that meets the limitation of “providing said public virtual IP address from said site switch to at least one load balancing controller.” Indeed, the figure on page 6 does not specifically show or describe any communication of public virtual IP addresses (such as communications represented by arrows) from the site switch in Hong Kong to the CGS.

Moreover, there is nothing in these cited passages of the White Paper involving the load balancing controller being enabled to “to update an address record to indicate said public virtual IP address as being associated with said site switch.” That is, nothing has been cited by the present Office Action from the White Paper that discloses, teaches, or suggests an address record and update thereof as recited in claim 34.

Accordingly, it is respectfully submitted that claim 34 is allowable over the White Paper.

As further evidence of the allowability of claim 34 over the subject matter disclosed in the White Paper, included herewith is an affidavit by Prajakta Joshi. As stated in the affidavit by Ms. Joshi, who is familiar with the subject matter described in the White Paper, the site switch did not communicate public virtual IP addresses to the CGS of the White Paper, prior to her invention. Instead for an architecture involving private virtual IP addresses mapped to public virtual IP addresses, a private virtual IP address (rather than a public virtual IP address) was communicated by the site switch to the CGS of the White Paper, thereby preventing the CGS from properly updating an address record to indicate the public virtual IP address (mapped to the private virtual IP address) as being associated with the site switch.

Accordingly, it is respectfully submitted that claim 34 is thus further allowable over the White Paper.

B. Independent claims 39, 43, 47, 51, and 56

Independent claims 43 and 51 recite limitations (using varying language) generally corresponding to those of claim 34, for a site switch. Independent claims 39, 47, and 56 are written from the point of view of a load balance switch, and recite for example, receiving a public virtual IP address and an address record that indicates the received public virtual IP address as being configured at the site switch.

For reasons similar to those explained above, it is respectfully submitted that claims 39, 43, 47, 51, and 56 are allowable as well.

C. Dependent claims 37, 40, 48, 54, and 57

Dependent claim 37 as amended herein recites, *inter alia*, that virtual public IP addresses received by the load balancing controller that “do not have indication in said address record as being associated with corresponding said site switches ... are excluded from having applied thereto any metric of a load balancing algorithm that is usable with virtual IP addresses.” It is respectfully submitted that the present Office Action has failed to cite any passage of the White Paper that meets these limitations.

For example in rejecting previously presented claim 37, page 5 of the present Office Action stated “See the paragraph in p. 2 [of the White Paper], discussing Scalability and Management, and refer to non-registered VIP addresses are added to the server farm, and transparent to the user.” This citation of the White Paper to reject claim 37 is traversed herein.

For example, the present Office Action has not explained how adding “non-registered VIP addresses” to a server farm and/or transparency to the user meets the limitations of “public virtual IP addresses ...that do not have indication in said address record as being associated with corresponding said site switches ... are excluded from having applied thereto any metric of a load balancing algorithm that is usable with virtual IP addresses.” Stated in another way, the present Office Action has not cited any portion of the White Paper that

speaks of virtual public IP addresses being excluded from having a metric applied thereto, in a manner required by claim 37.

Hence, claim 37 is allowable.

Dependent claims 40, 48, 54, and 57 recite limitations (using varying language) generally corresponding to claim 37, and are allowable for similar reasons.

D. Dependent claims 38, 42, 46, 50, 55, and 58

Dependent claim 38 as amended herein recites, *inter alia*, “an active bindings metric that prefers a virtual IP address, configured at respective said site switches, having a maximum number of active ones of said host servers bound to said preferred virtual IP address, rather than preferring another virtual IP address having a number of bound active ones of said host servers that is less than said maximum number.” In rejecting previously presented claim 38, page 5 of the present Office Action cited page 9 (the “High Availability” and “Maximum Scalability” subsections) of the White Paper. It is respectfully submitted that the present Office Action has not adequately identified any teachings in these cited passages that meet the specific limitations for an active bindings metric as recited in claim 38.

The passages on page 9 of the White Paper relied upon by the present Office Action are reproduced in full below:

“High Availability for Server and Application

ServerIron switches ensure service availability by offering switch, server, link, and session level redundancy. In the event of a server or application outage, ServerIron switches provide detection and sub-second fail-over to the next server that supports a like service. ServerIron switches provide active-standby redundancy or active-active redundancy capability that allows administrators to establish primary and secondary load balancing switches to support identical configurations parameters and provide 100 percent availability.

Maximum Scalability

TrafficWorks™ IronWare™ running on ServerIron simplifies network design by enabling IT managers to create a server farm, represented by a single IP address known as a virtual IP address (VIP). ServerIron acts as a single server for incoming Web traffic, controlling, monitoring and directing client requests to the most appropriate real server in a server farm. ServerIron's firewall load balancing (FWLB) eliminates firewall bottlenecks by distributing load across multiple firewalls. FWLB is Check Point certified, and can load balance up to 32 firewalls for a scalable and highly available Web-site deployment."

As clearly evident from the above-quoted passages from the White Paper, the "active bindings metric" as recited in claim 38 is not disclosed, taught, or suggested by these passages.

For example, the "High Availability" subsection is speaking of redundancy capability in which a "primary" load balancing switch is supported by a redundant "secondary" load balancing switch, in the event that the primary load balancing switch experiences a failure. It is respectfully submitted that the present Office Action has not explained how such load balancing switch "redundancy" teachings of the White Paper meet the limitations of claim 38 for an active bindings metric that prefers a virtual IP address (configured at a site switch) having a maximum number of active host servers bound to it, rather than another virtual IP address having a lesser number of active host servers. Stated in another way, the cited passage of the White Paper is speaking of providing redundancy between load balancing switches, whereas claim 38 is directed towards a load balancing controller that applies an active bindings metric that prefers a particular virtual IP address over another virtual IP address. The cited passage and claim 38 are thus addressing two different concepts.

As another example, the "Maximum Scalability" subsection is speaking of "firewall load balancing" (FWLB) in which traffic bottlenecks are eliminated by distributing the

traffic across multiple firewalls. Again, it is respectfully submitted that the present Office Action has not explained how such FWLB teachings of the White Paper meet the limitations of claim 38 for an active bindings metric that prefers a virtual IP address (configured at a site switch) having a maximum number of active host servers bound to it, rather than another virtual IP address having a lesser number of active host servers. Stated in another way, even if the firewalls of the White Paper might be *hypothetically* associated with virtual IP addresses, the White Paper makes no mention of preferring one virtual IP address over another—the White Paper simply says that load is “distributed across multiple firewalls” and provides no further detail as to any metrics that may be used to decide how such load is distributed.

Accordingly, it is respectfully submitted that claim 38 is allowable.

Claims 42, 46, 50, 55, and 58 recite limitations (using varying language) generally corresponding to claim 38, and are allowable for similar reasons.

E. Other claim amendments

Various other amendments are made to the claims as shown to provide appropriate antecedent basis, to make the language between and within the claims consistent given the amendments to the independent claims, to remove redundant and/or unnecessary limitations, to make changes of a typographical/grammatical nature, to more precisely recite the subject matter contained in the claims, and/or to otherwise place such claims in better form.

V. Conclusion

If there are any informalities or questions that can be addressed via telephone, the Examiner is encouraged to contact the attorney of record (Dennis M. de Guzman) at (206) 622-4900.

The Director is authorized to charge any additional fees due by way of this Amendment, or credit any overpayment, to our Deposit Account No. 19-1090.

Application No. 10/674,627
Reply to Office Action dated July 30, 2007

All of the claims remaining in the application are believed to be allowable.
Favorable consideration and a Notice of Allowance are earnestly solicited.

Respectfully submitted,
SEED Intellectual Property Law Group PLLC

/Dennis M. de Guzman/

Dennis M. de Guzman
Registration No. 41,702

DMD:wt

Enclosures:

Second Redlined Substitute Specification

Second Substitute Specification

Affidavit of Prajakta S. Joshi

701 Fifth Avenue, Suite 5400

Seattle, Washington 98104-7092

Phone: (206) 622-4900

Fax: (206) 682-6031

1017783

GLOBAL SERVER LOAD BALANCING SUPPORT FOR PRIVATE VIP ADDRESSES

BACKGROUND OF THE INVENTION

Field of the Invention

5 This disclosure relates generally to load balancing among servers. More particularly but not exclusively, the present disclosure relates to providing network components with capability to detect mapping between public and private addresses and to provide the public addresses for use in load balancing.

Description of the Related Art

10 Under the Transmission Control Protocol/Internet Protocol (TCP/IP), when a client provides a symbolic name (a Uniform Resource Locator or URL) to request access to an application program or another type of resource, the host name portion of the URL needs to be resolved into an IP address of a server for that application program or resource. For example, the URL (e.g.,
15 ~~http://www.foundrynet.com/index.htm~~ http colon double-slash dub-dub-dub dot foundrynet dotcom slash index dot htm) includes a host name portion ~~www.foundrynet.com~~ dub-dub-dub dot foundrynet dotcom that needs to be resolved into an IP address. The host name portion is first provided by the client to a local name resolver, which then queries a local Domain Name System (DNS)
20 server to obtain a corresponding IP address. If a corresponding IP address is not locally cached at the time of the query, or if the time-to-live (TTL) of a corresponding IP address cached locally has expired, the DNS server then acts as a resolver and dispatches a recursive query to another DNS server. This process
25 is repeated until an authoritative DNS server for the domain (e.g., ~~foundrynet.com~~ foundrynet dotcom, in this example) is reached. The authoritative DNS server

returns one or more IP addresses, each corresponding to an address at which a server hosting the application ("host server") under the host name can be reached. These IP addresses are propagated back via the local DNS server to the original resolver. The application at the client then uses one of the IP addresses to

5 establish a TCP connection with the corresponding host server. Each DNS server caches the list of IP addresses received from the authoritative DNS server for responding to future queries regarding the same host name, until the TTL of the IP addresses expires.

To provide some load sharing among the host servers, global server

10 load balancing (GSLB) switches are sometimes used as proxies for authoritative DNS servers, together with one or more site switches each associated with one or more host servers. Each site switch provides the GSLB switch with current site-specific information ("metrics") regarding access conditions to the host servers associated with the site switches. The GSLB switch then processes the addresses

15 returned by the DNS server using the metrics compiled from the site switches and provides an ordered address list having the optimum address for access listed at the top. An example of a GSLB system and description of associated metrics are disclosed in U.S. Application Serial No. 10/376,903, entitled "GLOBAL SERVER LOAD BALANCING," filed February 28, 2003, assigned to the same assignee as

20 the present application, and which is incorporated herein by reference in its entirety.

An increasingly common feature of networks with internal and external connections is the mapping of private (internal) server addresses to public (external) addresses via a mapping device, such a firewall or Network Address

25 Translation (NAT) device. Another frequent characteristic of such networks is the use of virtual IP addresses (VIPs) in addition to real server addresses. A VIP can have either or both a private address and a public address. The authoritative DNS server for which a GSLB switch is being used as a proxy for the specified domains is typically configured with the public addresses for these domains, so that the

GSLB switch can reorder these public addresses returned in the authoritative DNS server reply as part of the GSLB algorithm when a client requests access to any of the specified domains. In addition to having a GSLB switch communicate directly with site switches to obtain metrics information from the site switches, the GSLB switch also receives from the site switches a list of VIPs configured on the site switches. If these VIPs are private IP addresses mapped to public IP addresses by a device such as a firewall or NAT device, then the site switch is unaware of the mapping and only communicates the private VIP addresses to the GSLB switch. However, since the authoritative DNS server is configured with the public addresses rather than with the private addresses, the public VIP addresses received in the DNS reply do not match the private VIP address on the GSLB switch and are treated as real addresses by the GSLB switch rather than as virtual addresses. Since most of the metrics are applicable only to virtual addresses and not to real addresses, the GSLB switch cannot apply many of the metrics to the received private address, thereby reducing the overall efficiency or accuracy of the load balancing system.

As a further elaboration, a VIP having a private IP address is configured on a site switch. The site switch would know the private IP address associated with that VIP, but would not know the public IP address mapped to that private IP address by a mapping device (such as a firewall device). As a result, the site switch would communicate only the private IP address (and its associated metrics information) rather than the public IP address to the peer GSLB switch. Meanwhile, the authoritative DNS server (for which the peer GSLB switch is serving as a proxy and for which the GSLB switch is handling load balancing for the site having the VIP) has been configured with only the public IP address for the VIP for that site. Accordingly, when the GSLB switch receives the DNS reply from the authoritative DNS server, the GSLB switch would not recognize the public IP address in the DNS reply as being a VIP at that site, since the GSLB switch is only aware of the private IP address of the VIP received from the site switch. The

GSLB switch therefore treats the received public IP address as a real address, since the private IP address is different from the public IP address in the DNS reply being reordered by the GSLB switch. Accordingly, the GSLB switch would not apply (or would incorrectly apply) some of the metrics, such as the active bindings
5 metric (where the best IP address is the VIP that has the maximum number of active real servers bound to it), which are usable only with virtual addresses. Had the GSLB switch been able to correctly identify the received address as being a VIP, the GSLB would have been able to apply the correct metric(s) for VIPs when reordering the reply from the authoritative DNS server for which it is serving as a
10 proxy.

BRIEF SUMMARY OF THE INVENTION

One aspect of the present invention provides a method that includes obtaining information related to a mapping between first and second addresses associated with a network resource. The method sends the mapping information
15 to a load balancing device to allow the load balancing device to load balance traffic to the network resource using at least one metric associated with the second address and the mapping information.

The present invention is better understood upon consideration of the detailed description of the embodiments below, in conjunction with the
20 accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Non-limiting and non-exhaustive embodiments of the present invention are described with reference to the following figures, wherein like reference numerals refer to like parts throughout the various views unless
25 otherwise specified.

Figure 1 illustrates a GSLB system in which an embodiment of the invention may be implemented.

Figure 2 illustrates a GSLB system according to one embodiment of the invention with a remote load balancing arrangement.

Figure 3 illustrates a GSLB system according to another embodiment of the invention with a combination of remote and local load balancing
5 arrangement.

Figure 4 illustrates a GSLB system according to yet another embodiment of the invention with a combination of remote and local load balancing arrangement.

Figure 5 is a flowchart of a GSLB process to provide public
10 addresses according to an embodiment of the invention for a remote load balancing arrangement.

Figure 6 is a flowchart of a GSLB process to provide public addresses for only remote load balancing in an arrangement having both remote and local load balancing according to an embodiment of the invention.

15 Figure 7 is a flowchart of a GSLB process to provide public addresses for both remote and local load balancing according to an embodiment of the invention.

DETAILED DESCRIPTION

Embodiments of techniques to provide GSLB support for private VIPs
20 are described herein. In the following description, numerous specific details are given to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other methods, components, materials, etc. In other instances, well-known structures, materials, or operations are not shown or
25 described in detail to avoid obscuring aspects of the invention.

Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the

present invention. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

As an overview, one embodiment of the present invention provides GSLB support for private VIPs. According to this embodiment, an authoritative DNS server, for which a GSLB switch is serving as a proxy, is configured with public IP addresses for a domain that the GSLB switch load balances. A site switch for the GSLB switch is configured with one or more private address of a VIP of the domain, with the site switch providing metrics information to the GSLB switch as part of the load balancing process. A mapping device maps the private addresses of the VIP to the public IP addresses.

The site switch obtains the mapping information from the mapping device, thereby being able to identify the public IP address of the VIP. The site switch then communicates all of the VIPs, configured on the site switch, to the GSLB switch. Because the site switch has identified the public IP address that is mapped to the private IP address configured on the site switch, the communication of the VIPs from site switch to the GSLB switch includes the public IP address instead of the private IP address. The GSLB switch receives the public IP address of the VIP from the site switch. The GSLB switch updates an address list/records it maintains for the site switch with the public IP addresses of the VIP. As a result, when the GSLB switch reorders the DNS reply from the authoritative DNS server, the GSLB switch references the address list and correctly identifies the IP address in the reply as a VIP on the site switch, since the IP address configured for the domain on the authoritative DNS server as well as that learned by the GSLB switch from the site switch now is the public IP address of the VIP. The GSLB switch is now able to apply the appropriate VIP-related metrics accurately to reorder the DNS reply to send to a requesting client.

According to various implementations, the site switch can be configured “for peer only” or “for self and peer.” With the “for peer only” configuration, a GSLB controller on the site switch continues to use private IP addresses if the site switch also performs GSLB for the local site. With the “for self and peer” configuration, the site switch communicates the public IP addresses to a peer GSLB switch as well as to a local GSLB controller if the site switch is also functioning as a GSLB switch, thereby allowing the local GSLB controller of the site switch to accurately apply VIP-related metrics to load balance traffic.

Figure 1 shows one example of a global server load balancing arrangement in which an embodiment of the invention may be implemented. As shown in Figure 1, a remote (peer) GLSB switch 100 is connected to an Internet 104 and acts as proxy to an authoritative DNS server 102 for a network represented by the domain “~~foo.com~~foo dotcom” (for example). While the DNS server 102 provides the actual DNS service for the domain, the IP address known to the rest of the Internet 104 for the authoritative DNS server 102 is a VIP address configured on the GSLB switch 100. The DNS server 102 is also configured with the IP addresses for the domain ~~foo.com~~foo dotcom, and the GSLB switch 100 is configured as a proxy for the domain ~~foo.com~~foo dotcom. The GSLB switch 100 forwards client queries to the DNS server 102, and reorders the IP address list received from the authoritative DNS server 102 and sends the reordered IP address list in response to queries from clients requesting access to ~~foo.com~~foo dotcom.

The network represented by the domain name ~~foo.com~~foo dotcom has two components, for the purpose of describing an embodiment of this invention, in addition to other sub-parts. These components are a mapping device 106 and at least one site switch 108 (or other network device, such as a router). The mapping device 106 translates internal (private) addresses of real and virtual servers on the network to external (public) addresses. NAT or firewall devices are typical examples of such mapping devices 106.

The site switch 108 is coupled to an internal side of the mapping device 106. In addition to other tasks, the site switch 108 collects information about real and/or virtual servers 110 on the network and communicates with the GSLB switch 100. In particular, the site switch 108 has one or more VIPs
5 configured on it, and communicates to the GSLB switch 100 that it has these VIPs via a protocol exchange. This protocol exchange is also used to communicate VIP-related metrics information collected by the site switch 108 to the GSLB switch 100.

In a global server load balancing application, the GSLB switch 100,
10 acting as a proxy to the authoritative DNS server 102, receives a query from a client on the Internet 104 in the form of a URL that requests access to the domain ~~foo.com~~ foo dotcom, for example. The authoritative DNS server 102 provides a list of addresses to the GSLB switch 100 that corresponds to the domain-~~foo.com~~ foo dotcom. The GSLB switch 100 also gets metrics information along with a list of
15 VIPs configured on the site switch 108 from the site switch 108. Using the metrics information, the GSLB switch 100 reorders the list of addresses received from the authoritative DNS server 102 to place the optimum address at the top. For purposes of brevity, details of global server load balancing and performance metrics for load balancing will not be described in further detail herein, and instead
20 are disclosed in U.S. Application Serial No. 10/305,823, entitled "DISTRIBUTED HEALTH CHECK FOR GLOBAL SERVER LOAD BALANCING", filed November 27, 2002; U.S. Application Serial No. 10/376,903, entitled "GLOBAL SERVER LOAD BALANCING", filed February 28, 2003; and in U.S. Application Serial No. 10/211,822, entitled "STATISTICAL TRACKING FOR GLOBAL SERVER LOAD
25 BALANCING", filed August 1, 2002. All applications are assigned to the same assignee as the present application and incorporated herein by reference in their entirety.

In accordance with embodiments of the invention that will be described further below, the site switch 108 obtains mapping information between

public and private IP addresses on the network from the mapping device 106 and then communicates the VIPs configured on the site switch 108 to the GSLB switch 100, with the communication including the public IP address of the VIP rather than its private IP address. The GSLB switch 100 can then responsively update its address records 112 (e.g., a VIP list that the GSLB switch 100 maintains for the site switch 108) with the public IP address of the VIP. This public IP address of the VIP is also configured for the domain-~~foo.com~~ foo dotcom on the DNS server 102. The DNS server 102 returns a list of IP addresses, also containing the public IP address of the VIP, to the GSLB switch 100. The GSLB switch 100 refers to its address records 112 and correctly identifies the public IP address as a VIP on the site switch 108. Thus, the GSLB switch 100 can reorder the list of addresses received from the authoritative DNS server 102 based on the VIP-related metrics information and/or other metrics information provided by the site switch 108.

Examples of suitable topographies for load balancing with private VIP support include, but are not limited to, the following three arrangements in Figures 2-4. Figure 2 is a block diagram of a remote GSLB arrangement according to an embodiment of the invention. In this arrangement, a site switch 208 has no associated local GSLB sites and does not function as a GSLB switch itself (e.g., the site switch 208 has no local GSLB controller or metric collector). The site switch 208 operates as the remote site switch for a GSLB switch 200 (which load balances traffic to a network 224), and the GSLB switch 200 has no local site configured. Servers 210 have real IP addresses, and can have private VIP addresses configured on the site switch 208.

The network 224 can have a variety of mapping device arrangements. No mapping device, a mapping device integrated with the site switch 208, or an external mapping device connected to the site switch 208 are some of the examples. A mapping device 206 is shown in Figure 2, which may be a firewall, NAT, or other device that maps or otherwise allocates public IP addresses to the private IP addresses configured on the site switch 208. The

mapping information can be stored in one or more tables 220 or other data structure accessible by the site switch 208 in one embodiment.

There are several techniques that may be employed to allow the site switch 208 to obtain the mapping between the private and public IP addresses. In one embodiment, the mapping can be obtained via user configuration information. In this embodiment, a user can explicitly configure (such as via programming) the site switch 208 with the particulars of the mapping between the public IP addresses and the private IP addresses. In another embodiment, the mapping device 206 is integrated with the site switch 208, and the site switch 208 can directly obtain the mapping information from the table 220 (internal) or other entity that is maintained with the allocation of public IP addresses to private IP addresses. In yet another embodiment, the site switch 208 can obtain the mapping information through a message communication 226 with the mapping device 206, if the mapping device 206 is external to the site switch 208. The message communication 226 can be unidirectional or bi-directional movement of data between the site switch 208 and mapping device 206.

Then, the site switch 208 provides the obtained public IP addresses for the VIPs configured on it to the GSLB switch 200 that handles the remote load balancing for that particular network 224. According to one embodiment, the public IP addresses are provided to the GSLB switch 200 via a protocol message communication 222, along with related metrics information, instead of the private IP addresses. Alternatively or in addition in another embodiment, the information provided via the message communication 222 includes information indicative of the mapping between the public and private IP addresses, rather than solely the public IP addresses.

A specific example is now provided with regards to operation of the arrangement of Figure 2. Also provided to assist in explaining the operation is Figure 5, which is a flowchart of a GSLB process to provide public addresses according to an embodiment of the invention for the remote load balancing

arrangement of Figure 2. At least some components of the flowcharts depicted herein may be embodied in software or other machine-readable instructions stored on one or more machine-readable media. Such machine-readable media may be at a site switch, at a remote GSLB switch, or at other locations or combinations thereof. The various operations represented in the flowchart need not necessarily occur in the exact order depicted and some operations can be eliminated, added, or combined.

Initially in this example, the user configures a VIP 192.168.10.1 on the site switch 208. The VIP address 192.168.10.1 is a private IP address. The IP addresses (public) for a domain ~~www.gslb.com~~ dub-dub-dub dot gslb dot com are 207.95.55.23 and 253.72.96.55, and which are configured on the DNS server 202. The GSLB switch 200 is serving as a proxy to the DNS server 202 and is providing GSLB for the domain ~~www.gslb.com~~ dub-dub-dub dot gslb dot com.

The mapping device 206 maps the private IP address 192.168.10.1 of the VIP to one of the public IP addresses 207.95.55.23 (for example). The operations of Figure 5 begin with examples of different ways of providing mapping information to the site switch 208. Mapping information between the public and private IP addresses can be obtained by the site switch 208 via user configuration at a block 500, via access to internal allocation tables of an integrated mapping device at a block 502, via message communication with an external mapping device at a block 504, or via some other technique. Should the mapping information change at any time, the site switch 208 can re-learn the mapping. In addition to the mapping information (e.g., the public IP addresses), the site switch 208 also collects related metrics information at a block 506.

The site switch 208 is configured “for peer only” in this example, since the site switch 208 does not function as a GSLB controller/collector and does not have a local GSLB site configured. Therefore, the site switch 208 will be sending public IP addresses to the peer GSLB switch 200 only, to allow the peer GSLB switch 200 to perform load balancing accurately for the domain

~~www.gslb.com~~ dub-dub-dub dot gslb dotcom, rather than also sending public IP addresses to its internal GSLB components (which it does not have or are not enabled). The site switch 208, having obtained the mapping and metrics information, transmits that information to the GSLB switch 200 at a block 508.

- 5 More specifically, the site switch 208 notifies the GSLB switch 200 that it has a VIP 207.95.55.23 configured on it. The GSLB switch 200 maintains a list (e.g., address records 212) of VIPs for each site switch, and at a block 510, updates the address 207.95.55.23 (to indicate that this address is a VIP) in the VIP list maintained for the site switch 208.

- 10 A client makes a query to the GSLB switch 200 requesting access to ~~www.gslb.com~~ dub-dub-dub dot gslb dotcom at a block 512, with the IP addresses configured on the authoritative DNS server 202 for the domain ~~www.gslb.com~~ dub-dub-dub dot gslb dotcom being 207.95.55.23 and 253.72.96.55. The GSLB switch 200 forwards the request to the DNS server 202
- 15 at a block 511, and the DNS server 202 sends an address list associated with the domain to the GSLB switch at a block 515 (i.e., the addresses 207.95.55.23 and 253.72.96.55 in this example). The GSLB site switch 200 refers to the address records 212 at a block 513 and can now correctly identify that 207.95.55.23 is a VIP on the site switch 208, since the GSLB switch 200 now has this IP address in
- 20 the VIP list maintained for the site switch 208. The GSLB site switch 200 then performs GSLB on these IP address using the applicable metrics and selects the best IP address from among the addresses 207.95.55.23 and 253.72.96.55 at a block 514. Information reported by the site switch 208 can be used by the GSLB switch 200 for the VIP 207.95.55.23 during this selection process at the block 514,
- 25 since the IP address 207.95.55.23 has been correctly identified as a VIP on the site switch 208. The final operation in the flowchart is the transmission of the ordered address list to the inquiring client at a block 516.

In the load balancing arrangement depicted in Figure 3, both remote and local load balancing are performed, except that the arrangement is configured

“for peer only,” where public IP addresses are communicated only to a remote peer GSLB switch 300 from a site switch 308, which performs its own (“self”) local GSLB but does not communicate public IP addresses to its own internal GSLB components.

5 To further elaborate, the site switch 308 performs GSLB for one or more associated local sites 312 and remote sites (if any), and in addition is the site switch for the GSLB switch 300 that load balances traffic to a site 314 having host servers 310 coupled to the site switch 308. The site switch 308 is configured with the private VIP addresses to which the servers 310 of the site 314 are bound and
10 obtains at 324 mapping information from a mapping device 306, which maps these private VIP addresses to public IP addresses. The public IP addresses are obtained from the mapping device 306 using techniques previously described above, and as before with reference to Figures 2 and 5, the public IP addresses are communicated (along with related metrics information) by the site switch 308 to
15 the GSLB switch 300 via a protocol message communication 316, so that the GSLB switch 300 can update its VIP list with the public IP address of the VIP.

 The site switch 308 is also configured with the private IP addresses associated with the local site 312. However, since the site switch 308 is configured “for peer only,” the site switch 308 does not send any public IP
20 addresses associated with the local site 312 to the internal GSLB components 318 (e.g., a local GSLB controller or metric collector) integrated within the site switch 308. Rather, the site switch 308 sends the private VIP addresses configured on it to the internal GSLB components 318.

 An example description of the operation of the arrangement of Figure
25 3 will now be provided in conjunction with a flowchart of Figure 6. The GSLB switch 300 provides GSLB for a domain-~~www.gslb1.com~~ dub-dub-dub dot gslb1 dot com. The public IP addresses configured for ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dot com on an authoritative DNS server 302 are 207.95.55.23 and 253.72.96.55. As explained above, the GSLB switch 300 does not have any local

site configured, and has one remote site 314 with the site switch 308 (e.g., the site switch 308 is a site switch for the GSLB switch 300).

In addition to being a site switch for the GSLB switch 300, the site switch 308 itself is also a GSLB switch that provides GSLB for a domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom at the local site 312. The IP addresses configured for the domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom on an authoritative DNS server 322 are 192.168.10.1 and 192.168.72.1, which are private IP addresses. Accordingly, the site switch 308 provides GSLB for the domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom and is the local site switch used by its internal GSLB components 318. In addition, the site switch 308 operates as the remote site switch for the GSLB switch 300, which provides GSLB for the domain ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom.

When the “for peer only” configuration is completed on the site switch 308, the site switch 308 will do the following:

1) Since the site switch 308 is a remote site switch for the GSLB switch 300, the site switch 308 will communicate the VIPs configured on it via the message communication 316 to the GSLB switch 300. In particular, the site switch 308 obtains the mapping information including the public IP addresses for the site 314 from the mapping device 306 at a block 600, and will notify the GSLB switch 300 that the site switch 308 has a VIP 207.95.55.23 configured on it at a block 602. The GSLB switch 300 maintains address records 320 of VIPs for each site switch and at a block 604, updates the VIP address as 207.95.55.23 in the VIP records maintained for the site switch 308.

2) In addition, the site switch 308 is also a local site for the site switch 308 operating as a GSLB switch. Therefore, the site switch 308 will communicate the VIPs configured on it to its internal GSLB components 318 (along with metrics information) at a block 606. However, since the user has configured the “for peer only” option for one of the private VIP addresses (say 192.168.10.1, for instance), the local site switch 308 will notify its internal GSLB components 318 that it has a

private VIP 192.168.10.1 configured on it--note that it does not communicate the public IP address of the VIP to the internal GSLB components 318.

5 A client makes a query for ~~www.foo.com~~ dub-dub-dub dot foo dotcom, which the site switch 308 receives as a request to access the site 312 at a block 608, and forwards the request to the DNS server 322 at a block 609. The IP addresses configured for the domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom on the DNS server 322 are 192.168.10.1 and 192.168.72.1. The DNS server 322 sends the list of addresses containing 192.168.10.1 and 192.168.72.1 to the internal GSLB components 318. The internal GSLB components 318 refer to its address records 326 at a block 619 and correctly identify that 192.168.10.1 is a VIP configured on the local site switch 308, since the internal GSLB components 318 has this IP address in the VIP list it maintained for the site switch 312. The internal GSLB components 318 then perform GSLB on these IP addresses using the appropriate metrics and selects the best IP address for the client at a block 610. Metric information reported by the site switch 308 can be used by the internal GSLB components 318 for the VIP 192.168.10.1 during the selection process at the block 610. The prioritized list of addresses is sent to the requesting client at a block 612.

20 If a client makes a query for ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom to the GSLB switch 300 to request access to the site 314 at a block 614, the IP addresses configured for the domain ~~www.gslb.com~~ dub-dub-dub dot gslb dotcom on the DNS server 302 are 207.95.55.23 and 253.72.96.55. The GSLB switch 300 sends this request to the DNS server 302 at a block 615, and receives a list of addresses containing the addresses 207.95.55.23 and 253.72.96.55 from the DNS server 302 at a block 617. The GSLB switch 300 refers to its address records 320 at the block 619, and correctly identifies that 207.95.55.23 is a VIP on the site switch 308, since the GSLB switch 300 has this IP address in the VIP list maintained for the site switch 308. The GSLB switch 300 then performs GSLB on these IP addresses using the appropriate metrics and selects the best IP address

for the client at the block 610. Information reported by the site switch 308 can be used by the GSLB switch 300 for the VIP 207.95.55.23 during the selection process at the block 610, since the IP address 207.95.55.23 has been correctly identified as a VIP on the site switch 308.

5 Under another remote and local combination load balancing aspect of an embodiment of the invention shown in Figure 4, the site switch 308 can be configured “for self and peer,” wherein public IP address are provided by the site switch 308 to both the peer GSLB switch 300 and the internal GSLB components 318 integrated within the site switch 308. As with the embodiment of Figure 3, the
10 remote GSLB switch 300 receives metric information regarding servers 310 from the local site switch 308 and performs the load balancing for the site 314. The local site switch 308 also acts as an independent GSLB switch for the local site 312 and handles its load balancing based on the public IP addresses and metrics information received by its internal GSLB components 318 from other internal
15 components 400 of the site switch 308. A requesting client would query the remote GSLB switch 300 for addresses of domains it is providing GSLB for. Similarly, the requesting client would query the GSLB controller at the site switch 308 for addresses for domains the site switch 308 is providing GSLB for.

Another example is now provided to illustrate operation of the
20 arrangement of Figure 4 in connection with the “for self and peer” configuration, which is also to be described in connection with the flowchart of Figure 7. The GSLB switch 300 provides GSLB for the domain ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom. The public IP addresses configured for ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom on the DNS server 302 are 207.95.55.23 and
25 253.72.96.55. The GSLB switch 300 does not have any local site configured, and uses the site switch 308 as its remote site switch.

Additionally, the site switch 308 is also a GSLB switch providing
GSLB for the domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom, with the site switch 308 itself functioning as the site switch for its internal GSLB components

318. The public IP addresses configured for the domain ~~www.foo.com~~
dub-dub-dub dot foo dotcom on the DNS server 322 are 207.95.55.23 and
245.20.72.1, with the private IP address 192.168.10.1 being mapped to the public
IP address 207.95.55.23.

5 When the “for self and peer” configuration is performed on the site
switch 308, the site switch 308 will perform the following:

1) Since the site switch 308 is a remote site switch for the GSLB
switch 300, the site switch 308 will communicate the VIPs configured on the site
switch 308 via the message communication 316 to the GSLB switch 300. In
10 particular, the site switch 308 obtains (at 324 in Figure 4) the mapping information
including the public IP addresses for the site 314 from the mapping device 306 at a
block 700, and will notify the GSLB switch 300 that the site switch 308 has a VIP
207.95.55.23 configured on it at a block 702. The site switch 308 also sends
metrics information to the GSLB switch 300. The GSLB switch 300 maintains
15 address records 320 of VIPs for each site switch, and at a block 704, updates the
VIP address as 207.95.55.23 in the VIP records maintained for the site switch 308.

2) In addition, site switch 308 is also a local site for the internal GSLB
components 318. Therefore, components 400 will obtain mapping information
between private and public IP addresses associated with the site 312 at a block
20 706 and will communicate the VIPs configured on the site switch 308 to the
internal GSLB components 318 at a block 708. Since user has configured the “for
self and peer” option for the VIP 192.168.10.1, the internal components 400 will
notify the internal GSLB components 318 that the VIP 207.95.55.23 is configured
on the site switch 308 at the block 708--note that the public IP address of the VIP
25 is communicated to the internal GSLB components 318. If necessary, address
records 326 are updated by the internal GSLB components 318 to indicate that the
public IP address 207.95.55.23 corresponds to a VIP on the site switch 308.

The IP addresses configured for the domain ~~www.foo.com~~
dub-dub-dub dot foo dotcom on the DNS server 322 are 207.95.55.23 and

245.20.72.1. If a client makes a query to the site switch 308 at a block 712 for the domain ~~www.foo.com~~ dub-dub-dub dot foo dotcom, the site switch 308 forwards the request to the DNS server 322 at a block 713, and receives a list of addresses containing 207.95.55.23 and 245.20.72.1. The internal GSLB components 318
5 refers to the address records 326 at a block 719, and correctly identify that 207.95.55.23 is a VIP on the local site switch 308, since this public IP address is now kept in the VIP list (e.g., the address records 326) maintained for the site switch 308. The internal GSLB components 318 then perform GSLB on these IP addresses using the appropriate metrics and selects the best IP address for the
10 client at a block 714. Metric information reported by the site switch 308, or more particularly by the components 400, can be used by the internal GSLB components 318 for the VIP 207.95.55.23 during the selection process at the block 714.

The IP addresses configured for the domain ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom on the DNS server 302 are 207.95.55.23 and
15 dub-dub-dub dot gslb1 dotcom on the DNS server 302 are 207.95.55.23 and 253.72.96.55. If a client makes a query to the GSLB switch 300 for the domain ~~www.gslb1.com~~ dub-dub-dub dot gslb1 dotcom at a block 716, the GSLB switch 300 forwards the request to the DNS server 302 at a block 717, and receives a list of addresses containing 207.95.55.23 and 253.72.96.55. The GSLB switch 300
20 refers to the address records 320 at the block 719, and correctly identifies that 207.95.55.23 is a VIP on the site switch 318, since this IP address in the VIP list (e.g., the address records 320) maintained for the site switch 308. The GSLB switch 300 then performs GSLB on these IP addresses using the appropriate metrics and selects the best IP address for the client at the block 714. Information
25 reported by the site switch 308 can be used by the GSLB switch 300 for the VIP 207.95.55.23 during the selection process at the block 714, since the IP address 207.95.55.23 has been correctly identified as a VIP on the site switch 308. The prioritized list of addresses is sent to the requesting client at a block 718.

All of the above U.S. patents, U.S. patent application publications, U.S. patent applications, foreign patents, foreign patent applications and non-patent publications referred to in this specification and/or listed in the Application Data Sheet, are incorporated herein by reference, in their entirety.

5 The above description of illustrated embodiments of the invention, including what is described in the Abstract, is not intended to be exhaustive or to limit the invention to the precise forms disclosed. While specific embodiments of, and examples for, the invention are described herein for illustrative purposes, various equivalent modifications are possible within the scope of the invention and can be
10 made without deviating from the spirit and scope of the invention.

 These and other modifications can be made to the invention in light of the above detailed description. The terms used in the following claims should not be construed to limit the invention to the specific embodiments disclosed in the specification and the claims. Rather, the scope of the invention is to be
15 determined entirely by the following claims, which are to be construed in accordance with established doctrines of claim interpretation.

20

1017785_1.DOC